

SBBA: a Strongly-Budget-Balanced Double-Auction Mechanism

Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann

Bar-Ilan University, ramat-gan 5290002, Israel,
 {erelsgl,avinatanh,yaumann}@gmail.com

Abstract. In a seminal paper, McAfee (1992) presented the first dominant strategy truthful mechanism for double auction. His mechanism attains nearly optimal gain-from-trade when the market is sufficiently large. However, his mechanism may leave money on the table, since the price paid by the buyers may be higher than the price paid to the sellers. This money is included in the gain-from-trade and in some cases it accounts for almost all the gain-from-trade, leaving almost no gain-from-trade to the traders. We present SBBA: a variant of McAfee’s mechanism which is strongly budget-balanced. There is a single price, all money is exchanged between buyers and sellers and no money is left on the table. This means that all gain-from-trade is enjoyed by the traders. We generalize this variant to spatially-distributed markets with transit costs.

Keywords: mechanism design · double auction · budget balance · social welfare · gain from trade · spatially distributed market

1 Introduction

In the simplest *double auction* a single seller has a single item. The seller values the item for s , which is private information to the seller. A single buyer values the item for b , which is private to the buyer. If $b > s$, then trade can increase the utility for both traders; there is a potential *gain-from-trade* of $b - s$. However, there is no truthful, individually rational, budget-balanced mechanism that will perform the trade if-and-only-if it is beneficial to both traders. The reason is that it is impossible to determine a price truthfully. This is easy to see for a deterministic mechanism. If the mechanism chooses a price $p < b$, the seller is incentivized to bid $(p + b)/2$ to force the price up; similarly, if the mechanism chooses a price $p > s$, the buyer is incentivized to force the price down. The impossibility holds even when the valuations are drawn from a known prior distribution and even when the mechanism is allowed to randomize; see the classic papers of [13] and [16].

1.1 McAfee’s Trade-Reduction Mechanism

McAfee [12] showed how to circumvent this impossibility result when there are many sellers, seller i having private valuation s_i , and many buyers, buyer i having private valuation b_i . In McAfee’s double auction mechanism, each trader is

asked to give his valuation. The sellers are sorted in an ascending order according to their valuations $s_1 \leq s_2 \leq \dots \leq s_n$, and the buyers are sorted in a descending order $b_1 \geq b_2 \geq \dots \geq b_n$. Let k be the largest index such that $s_k \leq b_k$. The optimal gain-from-trade is attained by picking any price $p \in [s_k, b_k]$ and performing k deals in that price. But this scheme is not truthful. McAfee attains truthfulness by considering the following two cases:

(a) If there are at least $k + 1$ buyers and $k + 1$ sellers and the price $p_{k+1} := (b_{k+1} + s_{k+1})/2$ is in the range $[s_k, b_k]$, then p_{k+1} is set as the market price, allowing all k efficient deals to execute in that price.

(b) Otherwise, two prices are used, and $k - 1$ deals are done: all sellers with values s_1, \dots, s_{k-1} sell their item for s_k , and all buyers with values b_1, \dots, b_{k-1} buy an item for b_k . The mechanism performs a *trade reduction* by canceling a single deal, the deal between b_k and s_k , which is the least efficient of the k efficient deals. Hence, its gain-from-trade is $1 - 1/k$ of the maximum.

Crucially, the gain-from-trade approximated by McAfee's mechanism is the **total-gain-from-trade** - the gain-from-trade including the money left on the table due to the difference between the buyers' price and the sellers' price. Moreover, this money might include almost all the gain, so that the **market-gain-from-trade** - the gain enjoyed by the traders - might be near zero.

Example 1. There are k buyers and k sellers with the following valuations

- $s_i = 0$ and $b_i = B$ for all $i \in \{1, \dots, k - 1\}$.
- $s_k = \varepsilon$ and $b_k = B - \varepsilon$, where $0 < \varepsilon \ll B$.

The optimal market-gain-from-trade occurs when all sellers sell and all buyers buy, and it is: $k \cdot B - 2\varepsilon$.

Since there are no $k + 1$ -th buyer and seller, McAfee's mechanism sets the buy price at $B - \varepsilon$ and the sell price at ε . The trade includes $k - 1$ buyers and $k - 1$ sellers. The gain-from-trade in each deal is B , so the total-gain-from-trade is $(k - 1) \cdot B$, which is a very good approximation to the optimum for large k .

However, the net gain of each trader is ε , so the market-gain-from-trade is only $(k - 1) \cdot 2\varepsilon$; when $\varepsilon \rightarrow 0$, the market-gain-from-trade becomes arbitrarily small, and most gain-from-trade $(k - 1) \cdot (B - 2\varepsilon)$ remains on the table. \square

Money on the table can be desirable in some cases. E.g, the government may want to arrange a double-auction between commercial firms and collect the revenue. However, in other cases it may be considered unfair and drive traders away. For example, if traders in a stock-exchange notice that most gain-from-trade is taken by the operator, they may decide to switch to another operator.

The double-auction literature, e.g. [7], differentiates between mechanisms that are *weakly budget-balanced*, i.e, the auctioneer does not lose but may gain money, and *strongly budget-balanced*, i.e, the auctioneer does not lose nor gain any money. McAfee's mechanism is weakly budget-balanced.

1.2 Our Mechanism

In this paper we introduce SBBA - a Strongly-Budget-Balanced double-Auction mechanism. SBBA attains strong budget-balance by setting a *single* trade price

for all traders, in all cases. This may lead to excess supply; to handle the excess supply, a lottery is done between the sellers. At most one seller, selected at random, is excluded from trade. Hence, the expected total-gain-from-trade of SBBA is the same as McAfee's - $1 - 1/k$.

A disadvantage of SBBA is that its approximation holds only in expectation (taken over the randomization of the mechanism), while McAfee's approximation holds in the worst case. An advantage of SBBA is that it is strongly budget-balanced, so the market-gain-from-trade equals the total-gain-from-trade - all gain-from-trade is enjoyed by the traders. Besides these differences, SBBA has all the desirable properties of McAfee's mechanism:

- It is **ex-post individually-rational** - a trader never loses any value from participating in the market (a buyer is never forced to buy an item for more than its declared value; a seller is never forced to sell an item for less than its declared value; a trader who does not participate in the trade pays nothing).
- It is **ex-post dominant-strategy truthful**: for every trader, every vector of declarations by the other traders and every randomization, the trader's net value is always maximized by reporting his true value.
- It is **prior-free** - it does not assume or require any knowledge on the distribution of the traders' valuations. In other words, its approximation ratio is valid even for adversarial (worst-case) valuations.

Below we survey some related literature (Section 2). Then, we present the SBBA mechanism in the most basic double auction setting - a single market with single-unit buyers and single-unit sellers (Section 3). The idea of SBBA can be used in much more complex settings. To demonstrate its generality, we show how to use it in a *spatially-distributed market* - a collection of markets in different locations, with positive transit costs between markets (Section 4).

2 Related Work

2.1 Double auctions

VCG (Vickrey-Clarke-Grove) is a well-known mechanism that can be used in various settings, including double auction. It is truthful and attains the maximum gain-from-trade. Its main drawback is that it has budget deficit, which means that the auctioneer has to subsidize the market.

McAfee's Trade-Reduction mechanism was extended and generalized in many ways. Some of the extensions are surveyed below.

Babaioff et al extend McAfee's mechanism to handle spatially-distributed markets with transit costs [3] and supply chains [2,4,5], providing similar welfare guarantees. One variant, **Probabilistic Reduction** [2], achieves *ex-ante* budget balance by randomly selecting between the Vickrey-Clarke-Grove (VCG) mechanism and the Trade-Reduction mechanism. The probability is selected such that the deficit of the VCG exactly balances (in expectation) the surplus of McAfee. However, the probability depends on the distribution of the agents' valuations so the mechanism is not prior-free.

Lately, there has been a surge of interest in a more complicated market, namely double spectrum auctions, in which an auction is used to transfer spectrum from incumbent companies (e.g. TV stations) to modern companies (e.g. cellular operators). [20] adapted the Trade-Reduction mechanism to a double-spectrum-auction by creating groups of non-interfering buyers that can buy the same channel. [17] created an online variant of Trade-Reduction to handle the case in which new buyers arrive over time. [18] adapted Trade-Reduction to enable local markets, in which only some buyer-seller combinations are feasible. [10] adapted it to enable heterogeneous spectra. All these mechanisms are only weakly budget-balanced.

Some recent papers extend McAfee’s mechanism to settings with more than one item per trader. These are the **TAHES** mechanism of [10], which is multi-type single-unit; the **Secondary Market** mechanism of [19], which is single-type multi-unit; and the **Combinatorial Reallocation** mechanism of [7], which is multi-type multi-unit. Our mechanism is single-type single-unit; we leave to future work its extension to multi-type and multi-unit settings.

A different approach to double auctions is *random-sampling*. It was introduced by [6] under the assumptions that the traders’ valuations are random variables drawn from an unknown bounded-support distribution and there is a single item-type. We recently extended it to a prior-free setting with multiple item-types [15]. The idea is to divide the traders randomly to two half-markets, calculate an optimal price in one half and apply it to the other half, and vice versa. Since there is a single price in each market, the mechanism is strongly-budget-balanced. However, our analysis (for the single-type case) shows that the gain-from-trade is $1 - O(\sqrt{\ln k/k})$. While this still approaches 1 when the market is sufficiently large, the convergence rate is much slower than the $1 - 1/k$ guarantee of SBBA.

Recently, [9] presented a **two-sided sequential posted price mechanism (2SPM)**. Its main objective is to handle matroid constraints on the sets of buyers that can be served simultaneously. Like our mechanism, it is strongly budget-balanced. However, its approximation ratio is multiplicative - the ratio is between 4 to 16, depending on the setting. In particular, the approximation ratio does not approach 1 when the market is large.

The following table compares our work to some typical single-type single-unit double-auction mechanisms. In this table, an asterisk means "in expectation". TGFT means total-gain-from-trade and MGFT means market-gain-from-trade; they are identical for strongly-budget-balanced mechanisms.

| Mechanism | Prior-free | Budget | TGFT | MGFT |
|--|------------|----------|-------------------------|------|
| VCG | Yes | Deficit | 1 | 1 |
| Trade Reduction (McAfee 1992)[12] | Yes | Surplus | $1 - 1/k$ | 0 |
| Random Sampling (Baliga 2003)[6][15] | Yes | Balance | $1 - O(\sqrt{\ln k/k})$ | |
| Probabilistic Reduction (Babaioff 2009)[3] | No | Balance* | $1 - 1/k^*$ | |
| 2SPM (Colini 2016)[9] | No | Balance | $1/4$ to $1/16$ | |
| SBBA (this paper) | Yes | Balance | $1 - 1/k^*$ | |

2.2 Redistribution mechanisms

The problem illustrated by Example 1, where the money left on the table eats most of the welfare and leaves little welfare to the agents, happens in other domains besides double auctions. Several authors have suggested a two-step solution: in the first step, the original mechanism is executed and the budget-surplus is collected. In the second step, some of the surplus is re-distributed among the agents. This second step is called a *redistribution mechanism* and it should be carefully designed in order not to harm the truthfulness of the original mechanism. While it is not possible to redistribute the entire surplus in a truthful way, there are some truthful mechanisms that redistribute a large fraction of the surplus [8,11,1]. We take a different approach: we modify the original mechanism such that there is no budget-surplus at all, so no redistribution is needed and all social welfare remains with the agents.

3 The SBBA Mechanism


Order the buyers and sellers as in McAfee’s mechanism. In case of ties, impose an arbitrary order, e.g. lexicographic order of name.

Let k be the largest integer for which $s_k \leq b_k$ (or 0 if already $s_1 > b_1$). Call the first k sellers, the “cheap sellers”, and the first k buyers, the “expensive buyers”. As a convenience, if $k = n$, set $s_{k+1} = \infty$ and $b_{k+1} = 0$. Note that with this notation, we have that $s_i \leq b_i$ for $i \leq k$ and $s_i > b_i$ for $i > k$. The price is:

$$p := \min(s_{k+1}, b_k) .$$


There are two cases. We illustrate them below by plotting buyers’ valuations as balls and sellers’ valuations as squares (in all illustrations, $k = 3$).

Case 1: $s_{k+1} \leq b_k$ (note that by definition of k : $s_{k+1} > b_{k+1}$):



The price is $p = s_{k+1}$. All k expensive buyers and k cheap sellers trade in p .

Case 2: $s_{k+1} > b_k$:



The price is $p = b_k$. From the group of k cheap sellers, select $k - 1$ at random and let them trade with the $k - 1$ expensive buyers (excluding b_k).

Theorem 1. *The SBBA mechanism is prior-free (PF), individually-rational (IR), strongly-budget-balanced (SBB) and dominant-strategy truthful. Its expected market-gain-from-trade is at least $1 - 1/k$ of the optimum.*

Proof. The mechanism is PF by construction. It is IR since the trade is always between buyers whose value is above p and sellers whose value is below p . It is SBB since there is always a single price and all payments are between buyers and sellers. To analyze the gain-from-trade, note that in case 1, all k efficient deals are carried out and thus the maximum possible gain is achieved. In case 2, a single random deal is canceled, which implies an expected loss of $1/k$. Hence, in expectation, at most a fraction $1/k$ of the gain-from-trade is lost.

To prove truthfulness, we use a characterization of truthful single-parameter mechanisms from [14, chapter 9]. A mechanism is truthful iff:

- (a) the probability of an agent to win, given the bids of other agents, is a weakly monotonically increasing function of the agent's bid, and:
- (b) the price paid by a winning agent equals the *critical price* - the lowest value this agent has to bid in order to win, given the other agents' bids.

We prove that SBBA is truthful for the **buyers**. The winning probability of a buyer is either 0 or 1. Hence, it is sufficient to prove that a winning buyer never loses by raising the bid. Consider two cases of a winning buyer:

- The buyer is b_i for $i < k$.
- The buyer is b_k and $b_k \geq s_{k+1}$.

In both cases, if the bid is raised, b_k remains above s_{k+1} and the buyer's index may only decrease, so the buyer still wins.

The critical price when $b_k < s_{k+1}$ is b_k , since a trading buyer (one of the expensive $k - 1$ buyers) exits the trade by bidding below b_k and becoming the new b_k . The critical price when $b_k \geq s_{k+1}$ is s_{k+1} , since a trading buyer (one of the expensive k buyers) exits the trade by becoming b_k and bidding below s_{k+1} . In both cases, the price paid by the winning buyers is the critical price.

Finally, we prove that SBBA is truthful for the **sellers**. The sellers' ask-prices represent negative valuations, so monotonicity means that a seller's probability of participation should increase when the ask-price *decreases*. The winning probability of a seller is 0 when the seller is s_i for $i \geq k + 1$. When the seller is s_i for $i \leq k$, the winning probability is positive, and it does not depend on s_i itself but only on the relation between s_{k+1} and b_k :

- If $s_{k+1} \leq b_k$, the probability that s_i (for $i \leq k$) wins and pays is 1;
- If $s_{k+1} > b_k$, the probability that s_i (for $i \leq k$) wins and pays is $1 - 1/k$.

In each of these cases, decreasing the ask-price can only decrease the seller's index, so the winning probability remains the same.

The critical price when $s_{k+1} \leq b_k$ is s_{k+1} , since a trading seller (one of the k cheap sellers) exits the trade by asking above s_{k+1} and becoming the new s_{k+1} . This is indeed the price paid to a winning seller.

The critical price when $s_{k+1} > b_k$ is b_k , since a trading seller (one of the k cheap sellers) exits the trade by asking above b_k , which decreases k by 1 (the seller who increased his ask-price becomes s_k , but now the number of efficient deals is $k - 1$, so s_k is excluded from trade). Indeed, b_k is the price paid to a winning seller. Note that in this case both the winning and the price are realized with probability $1 - 1/k$.

□

3.1 Alternatives

The price set by our mechanism, $\min(b_k, s_{k+1})$, has an interesting economic interpretation: it is the highest price in a price-equilibrium (aka Walrasian equilibrium), i.e. the highest price in which the market can be cleared by balancing supply and demand. If the price is raised above b_k then the demand becomes less than k while the supply is still at least k ; if the price is raised above s_{k+1} then the supply becomes at least $k + 1$ while the demand is still at most k ; in both cases there is an excess supply.

It is possible to switch the role of buyers and sellers, splitting the cases by whether $b_{k+1} > s_k$. In this case, the price is $\max(s_k, b_{k+1})$ which is the *lowest* price-equilibrium. The alternative mechanism has the same properties of our original mechanism.

There are some other alternatives that come to mind, but are either not truthful or not efficient:

- If in Case 2, instead of using a lottery we select the trading sellers deterministically (e.g. taking the $k - 1$ cheaper sellers), the mechanism will not be truthful, since some agents who want to trade at the market price have an incentive to deviate from their true values in order to enter the trade.
- If we always set the price to s_{k+1} , in some cases this price might be higher than the valuations of all buyers so we might lose all gain-from-trade.
- If we always set the price to b_k , in some cases this price might be higher than the valuations of all sellers, the market will be flooded by inefficient sellers, and the expected gain-from-trade will be low.

4 Spatially Distributed Markets


A *spatially distributed market* is a collection of several markets, each of which is located in a different geographic location. It is possible to transport goods from one market to another for a fixed, positive transit cost, which may be different for each ordered pair of markets. Babaioff et al [3] extended McAfee’s mechanism to handle such markets. Similarly to McAfee’s mechanism, their mechanism has a budget surplus. Below we briefly present their mechanism and present a strongly budget-balanced variant of it.

4.1 Create the market-flow graph

Create a network-flow graph representing the market in the following way:


- Create a node for each market. Create a directed edge between each pair of markets, with infinite capacity and with cost equal to the transit cost between the two markets (which may be different in each direction).
- Create an additional *Agents node*, representing the buyers and sellers. For each seller in market i , create an edge FROM the Agents node TO Market i , with unit capacity and cost equal to the seller's ask-price. This edge represents the seller producing an item and sending it to the market. For each buyer in market i , create an edge TO the Agents node FROM Market i , with unit capacity and cost equal to MINUS the buyer's bid. This edge represents the buyer bringing an item from the market.

The following illustration shows a graph representing two markets, with a transit cost of 4 in each direction. Each solid arc represents an infinite-capacity edge. Each dashed arc represents several unit-capacity edges with different costs. The numeric labels are the costs.



4.2 Calculate a minimum-cost flow

Babai et al use a known polynomial-time algorithm for finding a flow with minimum cost in the market graph. Assuming all data is integral, the flow in every edge is also an integer number. In particular, the flow in every buyer/seller edge (with capacity 1) is either 0 or 1. Hence, a flow in the graph defines an allocation in which each trader trades if and only if the flow in the corresponding edge to/from the Agents node is 1. A minimum-cost flow corresponds to an optimal trade: the (negative) cost of the flow is minus the gain-from-trade. The following illustration shows the minimum-cost flow and the corresponding optimal trade in the above example market (non-trading agents are bracketed):



There are 6 efficient deals. The net cost is -100 so the gain-from-trade is 100.

4.3 Find commercial-relationship components

In the optimal flow, the markets can be partitioned to groups, such that all trade is within groups and no trade is between groups. Such groups are called *commercial-relationship components*; they are the connectivity components of a graph in which the nodes are the markets and there is an edge between markets trading in the optimal flow.

The optimal trade can be attained in a price-equilibrium, in which there is a single price in each market. In each component, the prices in the different markets are tied by the equilibrium conditions: if the price in market i is p_i , and there is positive trade from market i to market j , then we must have:

$$p_j = p_i + \text{Cost}[i, j] \quad (1)$$

since in equilibrium, the sellers in market i should be indifferent between selling in their local market for a net revenue of p_i , and selling in market j for a net revenue of $p_j - \text{Cost}[i, j]$. Therefore, in each component, setting the price in a single market uniquely determines the prices in other markets. Formally, for every two markets i, j in the same component there is a constant $\Delta_{i,j}$ such that, in any price-equilibrium, $p_j = p_i + \Delta_{i,j}$ ($\Delta_{i,j}$ can be calculated, for example, by calculating cheapest paths in the residual graph of the min-cost flow; see [3]).

In our running example, there is a single component. Here, $\Delta_{1,2} = \text{Cost}[1, 2] = 4$. The optimal trade can be attained in a price-equilibrium in which the price in Market 1 is p_1 and the price in Market 2 is $p_2 = p_1 + 4$. Any price-vector between $p_1 = 15, p_2 = 19$ and $p_1 = 17, p_2 = 21$ is an equilibrium price-vector.

4.4 Trade reduction

At this point, Babaioff et al calculate a *reduced residual graph* of the min-cost flow, remove a single cycle representing the least efficient deal, and determine

two prices in each market (buy price and sell price) based on distances in the reduced residual graph. This gives a truthful mechanism with a budget surplus.

Here our mechanism takes a different approach:

- In each commercial-relationship-component, virtually bring all traders to an arbitrary market in that component, e.g, Market i . Adjust their bids according to the price-equilibrium conditions: a bid b in Market j is translated to a bid $b - \Delta_{i,j}$ in Market i .
- Proceed as in the single-market situation (Section 3): order the buyers decreasingly and the sellers increasingly and find k - the total number of efficient deals in the component. Set the price in Market i to $p_i := \min(b_k, s_{k+1})$.
- Determine the prices in the other markets of the same component according to the equilibrium conditions: for every market j , set $p_j := p_i + \Delta_{i,j}$. Below we prove that all these prices are non-negative.
- If b_k is the price-setter, then exclude b_k and a random seller in the component from trading. Otherwise, allow all k efficient traders in the component to trade in their market prices.
- Ban any trade between different components.

In the above example, when all traders are brought to Market 1, we have the following valuations (here $\Delta_{1,2} = 4$; the adjusted valuations of traders brought from Market 2 are displayed in slanted digits):

- Sellers: $-2, 01, 05, 09, 13, 15, 17, 19, 23, 27$
- Buyers: $32, 28, 24, 20, 19, 18, 14, 12, 08, 04$

Here, the number of efficient deals $k = 6$. We have $b_k = 18$ and $s_{k+1} = 17$, so the price-setter is s_{k+1} , who is originally a seller in Market 2 (where his valuation is 21). The price-vector is $p_1 = 17, p_2 = 21$. All 6 efficient deals are performed. Another example, where the price-setter is b_k , is shown in Appendix A.

4.5 Analysis

Similarly to our single-market mechanism, the spatially-distributed-market mechanism is prior-free and strongly-budget-balanced, since there is a single price in each market. The reduction in trade is only at most a single deal *in each component*. Hence, the expected gain-from-trade in each component is $(1 - 1/k)$ of the optimum in that component, where k is the number of efficient deals in the component. When the components are large, this k may be much larger than the number of efficient deals in each market alone.

For truthfulness, we use the monotonicity characterization shown in Theorem 1.

First, we have to prove that a winning trader never loses by increasing/decreasing the bid/ask price. Indeed, if a trader is winning, then the edge from the Agents node to the trader is active in the min-cost flow. When the bid/ask increases/decreases, the cost of that edge decreases. Hence, the cost of the flow decreases, and it remains the min-cost flow. Thus, the partition of the graph to commercial relationship components does not change. Within each component, increasing/decreasing the bid/ask price weakly decreases the index of the trader in the ordering, so a winning buyer/seller is still among the first k buyers/sellers.

Next, we have to prove that each trader pays the critical price. For traders originally from Market i , the critical price is $p_i = \min(b_k, s_{k+1})$; this follows immediately from the proof of Theorem 1. Consider now a winning buyer from Market $j \neq i$ in the same component as Market i . If our buyer bids b , then the bid is translated to Market i as $b - \Delta_{i,j}$. The buyer exits the trade when $b - \Delta_{i,j} < p_i$. Hence, the critical price for our buyer is $p_i + \Delta_{i,j}$, which is exactly the price p_j paid by our buyer. Similar considerations are true for the sellers.

Finally, as promised, we prove that the prices determined by our mechanism are non-negative. Let Market j be a market with a smallest price in some component. Suppose by contradiction that $p_j < 0$. Since we proved that the mechanism is truthful, it implies that there are no active sellers in Market j (since a seller would prefer to lie than to sell in negative price). This means that there is no trade outgoing from Market j . Since there must be active sellers elsewhere in the component, there must be other markets in the component, and this means that there is trade incoming to Market j , say, from Market i . But this means that $p_i = p_j - \text{Cost}[i, j]$. Since we assume that transit costs are positive, this contradicts the minimality of p_j . \square

5 Future Work


1. Our strongly-budget-balanced mechanism is randomized. This is in contrast to the VCG and McAfee mechanisms, which are deterministic. This raises the following open question: is there a deterministic mechanism that is strongly-budget-balanced (in addition to being prior-free, truthful, individually-rational and approximately-optimal)?

2. Besides the spatially-distributed markets described in this paper, there are many other variants of McAfee's mechanism. An interesting line of future work is to survey these variants and see if and how they can be made strongly-budget-balanced.


Acknowledgments. This research was funded in part by the following institutions: The Doctoral Fellowships of Excellence Program at Bar-Ilan University, the Mordechai and Monique Katz Graduate Fellowship Program, and the Israel Science Fund grant 1083/13. We are grateful to Moshe Babaioff for his advice on the spatially-distributed-market mechanism.

A Another Example of a Spatially-Distributed Market

Consider a spatially-distributed market represented by the following graph:



Here is the optimal flow:



It can be attained in a price-equilibrium where the price-vector can be anywhere between $p_1 = 15, p_2 = 19$ and $p_1 = 16, p_2 = 20$ (with $p_2 = p_1 + 4$).

When all traders are brought to Market 1, we get:

- Sellers: 01, 05, 09, 11, 13, 15, 17, 18, 23, 27
- Buyers: 32, 28, 24, 20, 19, 16, 14, 12, 08, 04

The number of efficient deals $k = 6$. $b_k = 16$ and $s_{k+1} = 17$, so the price-setter is b_k (who is originally from Market 1). The price-vector is set to $p_1 = 16, p_2 = 20$. Buyer 16 from market 1 and one randomly-selected seller are excluded from trading. All in all, 5 out of 6 efficient deals are performed.

References

1. Apt, K., Conitzer, V., Guo, M., Markakis, E.: Welfare Undominated Groves Mechanisms. In: Papadimitriou, C., Zhang, S. (eds.) *Internet and Network Economics*, Lecture Notes in Computer Science, vol. 5385, pp. 426–437. Springer Berlin Heidelberg (2008)
2. Babaioff, M., Nisan, N.: Concurrent Auctions Across the Supply Chain. *Journal of Artificial Intelligence Research (JAIR)* 21, 595–629 (2004)
3. Babaioff, M., Nisan, N., Pavlov, E.: Mechanisms for a spatially distributed market. *Games and Economic Behavior* 66(2), 660–684 (Jul 2009)
4. Babaioff, M., Walsh, W.E.: Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. *Decision Support Systems* 39(1), 123–149 (Mar 2005)
5. Babaioff, M., Walsh, W.: Incentive Compatible Supply Chain Auctions. In: Chaidraa, B., Müller, J. (eds.) *Multiagent based Supply Chain Management*, vol. 28, pp. 315–350. Springer Berlin Heidelberg (2006)
6. Baliga, S., Vohra, R.: Market Research and Market Design. *Advances in Theoretical Economics* 3(1) (Jan 2003)
7. Blumrosen, L., Dobzinski, S.: Reallocation Mechanisms. In: *Proceedings of the Fifteenth ACM Conference on Economics and Computation*. p. 617. EC '14, ACM, New York, NY, USA (2014)
8. Cavallo, R.: Optimal Decision-making with Minimal Waste: Strategyproof Redistribution of VCG Payments. In: *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*. pp. 882–889. AAMAS '06, ACM, New York, NY, USA (2006)
9. Colini-Baldeschi, R., de Keijzer, B., Leonardi, S., Turchetta, S.: Approximately Efficient Double Auctions with Strong Budget Balance. In: *ACM Symposium on Discrete Algorithms* (Jan 2016)
10. Feng, X., Chen, Y., Zhang, J., Zhang, Q., Li, B.: TAHES: Truthful double Auction for Heterogeneous Spectrums. In: *INFOCOM, 2012 Proceedings IEEE*. pp. 3076–3080. IEEE (Mar 2012)
11. Guo, M., Conitzer, V.: Optimal-in-expectation Redistribution Mechanisms. In: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 2*. pp. 1047–1054. AAMAS '08, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2008)
12. McAfee, R.P.: A dominant strategy double auction. *Journal of Economic Theory* 56(2), 434–450 (Apr 1992)
13. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29(2), 265–281 (Apr 1983)
14. Nisan, N.: Introduction to Mechanism Design (For Computer Scientists). In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*, pp. 209–241 (2007)
15. Segal-Halevi, E. and Hassidim, A. and Aumann, Y.: MIDA: a Multi Item-type Double-Auction Mechanism. arXiv preprint <http://arxiv.org/abs/1604.06210> (2016)
16. Vickrey, W.: Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16(No. 1 (Mar., 1961)), 8–37 (1961)
17. Wang, S., Xu, P., Xu, X., Tang, S., Li, X., Liu, X.: TODA: Truthful Online Double Auction for Spectrum Allocation in Wireless Networks. In: *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on*. pp. 1–10. IEEE (Apr 2010)

18. Wang, W., Li, B., Liang, B.: District: Embracing local markets in truthful spectrum double auctions. In: Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2011 8th Annual IEEE Communications Society Conference on. pp. 521–529. IEEE (Jun 2011)
19. Xu, H., Jin, J., Li, B.: A Secondary Market for Spectrum. In: INFOCOM, 2010 Proceedings IEEE. pp. 1–5. IEEE (Mar 2010)
20. Yao, E., Lu, L., Jiang, W.: An efficient truthful double spectrum auction design for dynamic spectrum access. In: Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM), 2011 Sixth International ICST Conference on. pp. 181–185. IEEE (Jun 2011)